

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Shape and Animation Methods and Systems Using  
Examples**

Inventor(s):

Michael F. Cohen

Charles F. Rose, III

Peter-Pike J. Sloan

1 **TECHNICAL FIELD**

2 This invention relates to the fields of computer graphics and computer  
3 animation. More particularly, the invention concerns methods and systems for  
4 creating or blending animated forms and/or motions.

5  
6 **BACKGROUND**

7 Modern animation and modeling systems enable artists to create high-  
8 quality content, but provide limited support for interactive applications. Although  
9 complex forms and motions can be constructed either by hand or with motion or  
10 geometry capture technologies, once they are created, they are difficult to modify,  
11 particularly at runtime.

12 The magic of computer graphics as seen in many current movies and  
13 computer games comes at a cost. Creating geometric forms that are used to  
14 generate synthetic characters, and animating the characters to bring them to life  
15 requires either highly skilled artists and/or sophisticated capture devices. Both are  
16 expensive and, in the case of highly skilled artists, rare.

17 At this time, a variety of 3D scanning methodologies are available that can  
18 capture shapes that exist in the real world. Motion capture technologies are also  
19 capable of recording complex performances. One drawback of these capture  
20 systems is the cost that is incurred in their purchase and operation. Another  
21 limitation of capture devices is that they are restricted to recording shapes and  
22 motions that can be observed and that actually exist (e.g., such systems cannot  
23 capture the motion of a dinosaur). In contrast to the scanning devices, highly  
24 skilled artists have the advantage of being able to model complex existing shapes,  
25 as well as imaginary shapes and animations.

Both of these means of creating shapes and animations are limited in the same way. In particular, neither has a simple means of automatically modifying the shapes and animations once they have been created. It would be highly desirable for current systems to be able to automatically modify a shape or animation for two reasons. First, in scripted settings, such as films, automatic modification makes it easier to avoid redundancy. For instance, an entire colony of different looking ants could be automatically created from a few distinct ants. In addition, the ants' motion could automatically adapt to changes in their mood or to the terrain they are walking on. Second, in non-scripted interactive runtime settings, such as games, it is not possible to anticipate all shapes and animations that will be needed when the game is played. Thus, most games simply try to reuse the closest fitting model or motion for the situation. One solution to the above problems is to create more shapes and animations. Unfortunately, this approach is very costly.

Accordingly, this invention arose out of concerns associated with improving modern animation and modeling systems.

## **SUMMARY**

Methods and systems that are suitable for leveraging existing forms and motions to automatically generate variations are described. The methods and systems are particularly well suited for generating variations at run time or "on the fly". In the described embodiment, methodologies are presented for efficient runtime interpolation or blending between multiple forms or multiple motion segments. Interpolation provides a way to leverage artist-generated source

1 material. Radial basis functions provide mathematical support for the  
2 interpolation.

3 In the described embodiment, example forms or motions are provided. The  
4 forms or motions can be created by artists, or through geometry or motion capture  
5 technologies. The described system generates a continuous range of forms termed  
6 a "shape" or a continuous range of motions termed a "verb". Shape interpolation  
7 methodology can also be applied to articulated figures to create smoothly skinned  
8 figures that deform in natural ways. The runtime interpolation of the forms or  
9 motions runs fast enough to be used in interactive applications such as games.

### 10 11 **BRIEF DESCRIPTION OF THE DRAWINGS**

12 The file of this patent contains at least one drawing executed in color.  
13 Copies of this patent with color drawing(s) will be provided by the Patent and  
14 Trademark Office upon request and payment of the necessary fee.

15 Fig. 1 is a block diagram of an exemplary computer system that is suitable  
16 for use with the described embodiments.

17 Fig. 2 is a block diagram of an exemplary system in which the inventive  
18 methodologies can be employed.

19 Fig. 3 is a diagram that illustrates an exemplary abstract space and is useful  
20 in understanding some of the concepts behind the inventive embodiments.

21 Fig. 4 is a flow diagram that describes steps in a method in accordance with  
22 the described embodiments.

23 Fig. 5 is a flow diagram that describes steps in a method in accordance with  
24 the described embodiments.  
25

1 Fig. 6 is an illustration that is useful in understanding some of the concepts  
2 behind the inventive interpolation methods of the described embodiments.

3 Fig. 7 is a graph that illustrates linear and radial basis portions of a cardinal  
4 basis function in accordance with an example that is given.

5 Fig. 8 is a graph that illustrates cardinal basis functions and a scaled sum  
6 function for a particular degree of freedom in accordance with the Fig. 7 example.

7 Fig. 9 is a color picture that describes exploration of abstract space that can  
8 reveal problems involved with the described interpolation.

9 Fig. 10 is a color picture that describes pseudo-examples that can be used to  
10 re-parameterize abstract space to address the problems revealed and illustrated in  
11 Fig. 9.

12 Fig. 11 is a flow diagram that describes steps in a method in accordance  
13 with the described embodiment.

14 Fig. 12 is a color picture that illustrates simple transformation blending that  
15 leads to or exhibits shrinking about the joints.

16 Fig. 13 is a color picture that illustrates example forms that are warped into  
17 a cononical pose.

18 Fig. 14 is a color picture that illustrates naïve transform blending versus  
19 interpolated blending.

20 Fig. 15 is a flow diagram that describes steps in a method in accordance  
21 with the described embodiment.

22 Fig. 16 is a color picture that illustrates examples and pseudo-examples in  
23 accordance with the described embodiment.  
24  
25

1 Fig. 17 is a color picture that illustrates exemplary motions in accordance  
2 with the described embodiment. In the figure, health increases as  $y$  increases and  
3 slope of the surface ranges from sloping to left to right.

4 Fig. 18 is a color picture that illustrates exemplary motions in accordance  
5 with the described embodiment. In the figure, knowledge increases from left to  
6 right and happiness increases from bottom to top.

7 Fig. 19 is a color picture that provides more detail of the characters of Figs.  
8 17 and 18.

## 9 10 **DETAILED DESCRIPTION**

### 11 **Overview**

12 Modern animation and modeling systems enable artists to create high-  
13 quality content, but provide limited support for interactive applications. Although  
14 complex forms and motions can be constructed either by hand or with motion or  
15 geometry capture technologies, once they are created, they are difficult to modify,  
16 particularly at runtime.

17 Interpolation or blending provides a way to leverage artist-generated  
18 source material. Presented here are methodologies for efficient runtime  
19 interpolation between multiple forms or multiple motion segments. Radial basis  
20 functions provide key mathematical support for the interpolation. Once the  
21 illustrated and described system is provided with example forms and motions, it  
22 generates a continuous range of forms referred to as a "shape" or a continuous  
23 range of motions referred to as a verb.

24 Additionally, shape interpolation methodology is applied to articulated  
25 figures to create smoothly skinned figures that deform in natural ways. The

runtime interpolation of the forms or motions runs fast enough to be used in interactive applications such as games.

### **Exemplary Operating Environment**

Fig. 1 is a high level block diagram of an exemplary computer system 130 in which the described embodiments can be implemented. Various numbers of computers such as that shown can be used in the context of a distributed computing environment. These computers can be used to render graphics and process images in accordance with the description given below.

Computer 130 includes one or more processors or processing units 132, a system memory 134, and a bus 136 that couples various system components including the system memory 134 to processors 132. The bus 136 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. The system memory 134 includes read only memory (ROM) 138 and random access memory (RAM) 140. A basic input/output system (BIOS) 142, containing the basic routines that help to transfer information between elements within computer 130, such as during start-up, is stored in ROM 138.

Computer 130 further includes a hard disk drive 144 for reading from and writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and writing to a removable magnetic disk 148, and an optical disk drive 150 for reading from or writing to a removable optical disk 152 such as a CD ROM or other optical media. The hard disk drive 144, magnetic disk drive 146, and optical disk drive 150 are connected to the bus 136 by an SCSI interface 154 or some

1 other appropriate interface. The drives and their associated computer-readable  
2 media provide nonvolatile storage of computer-readable instructions, data  
3 structures, program modules and other data for computer 130. Although the  
4 exemplary environment described herein employs a hard disk, a removable  
5 magnetic disk 148 and a removable optical disk 152, it should be appreciated by  
6 those skilled in the art that other types of computer-readable media which can  
7 store data that is accessible by a computer, such as magnetic cassettes, flash  
8 memory cards, digital video disks, random access memories (RAMs), read only  
9 memories (ROMs), and the like, may also be used in the exemplary operating  
10 environment.

11 A number of program modules may be stored on the hard disk 144,  
12 magnetic disk 148, optical disk 152, ROM 138, or RAM 140, including an  
13 operating system 158, one or more application programs 160, other program  
14 modules 162, and program data 164. A user may enter commands and  
15 information into computer 130 through input devices such as a keyboard 166 and a  
16 pointing device 168. Other input devices (not shown) may include a microphone,  
17 joystick, game pad, satellite dish, scanner, or the like. These and other input  
18 devices are connected to the processing unit 132 through an interface 170 that is  
19 coupled to the bus 136. A monitor 172 or other type of display device is also  
20 connected to the bus 136 via an interface, such as a video adapter 174. In addition  
21 to the monitor, personal computers typically include other peripheral output  
22 devices (not shown) such as speakers and printers.

23 Computer 130 commonly operates in a networked environment using  
24 logical connections to one or more remote computers, such as a remote computer  
25 176. The remote computer 176 may be another personal computer, a server, a



1 router, a network PC, a peer device or other common network node, and typically  
2 includes many or all of the elements described above relative to computer 130,  
3 although only a memory storage device 178 has been illustrated in Fig. 1. The  
4 logical connections depicted in Fig. 1 include a local area network (LAN) 180 and  
5 a wide area network (WAN) 182. Such networking environments are  
6 commonplace in offices, enterprise-wide computer networks, intranets, and the  
7 Internet.

8 When used in a LAN networking environment, computer 130 is connected  
9 to the local network 180 through a network interface or adapter 184. When used  
10 in a WAN networking environment, computer 130 typically includes a modem 186  
11 or other means for establishing communications over the wide area network 182,  
12 such as the Internet. The modem 186, which may be internal or external, is  
13 connected to the bus 136 via a serial port interface 156. In a networked  
14 environment, program modules depicted relative to the personal computer 130, or  
15 portions thereof, may be stored in the remote memory storage device. It will be  
16 appreciated that the network connections shown are exemplary and other means of  
17 establishing a communications link between the computers may be used.

18 Generally, the data processors of computer 130 are programmed by means  
19 of instructions stored at different times in the various computer-readable storage  
20 media of the computer. Programs and operating systems are typically distributed,  
21 for example, on floppy disks or CD-ROMs. From there, they are installed or  
22 loaded into the secondary memory of a computer. At execution, they are loaded at  
23 least partially into the computer's primary electronic memory. The invention  
24 described herein includes these and other various types of computer-readable  
25 storage media when such media contain instructions or programs for implementing

1 the steps described below in conjunction with a microprocessor or other data  
2 processor. The invention also includes the computer itself when programmed  
3 according to the methods and techniques described below.

4 For purposes of illustration, programs and other executable program  
5 components such as the operating system are illustrated herein as discrete blocks,  
6 although it is recognized that such programs and components reside at various  
7 times in different storage components of the computer, and are executed by the  
8 data processor(s) of the computer.

### 9 10 **Exemplary System**

11 Fig. 2 shows an exemplary system 200 in which the inventive  
12 methodologies discussed below can be employed. In the discussion immediately  
13 below, the system's components are first identified followed by a discussion of  
14 their functionality.

### 15 16 **System Components**

17 In the illustrated example, system 200 comprises a designer component  
18 202, a shape/verb generator 204, and an exemplary application 206. One primary  
19 focus of system 200 is to provide a way, in an application, to have a very  
20 lightweight, fast runtime system that allows for modification of a shape or  
21 animation on the fly in the context of the application (which can typically  
22 comprise a game).

23 Designer component 202 comprises an external tools module that can  
24 include, among other components, modeling/animation systems 210 and  
25 shape/motion capture systems 212. The designer component also comprises

1 various shape and verb offline tools 214 that comprise an abstract space definition  
2 module 216 and an example position module 218 that positions examples in  
3 abstract space.

4 Application 206 comprises a shape/verb runtime system 220 that is  
5 embedded in an application, such as a game, that uses the system described below  
6 to drive the animation of a character. The inventive aspects that are about to be  
7 described pertain to the shape and verb offline tools 214, shape/verb generator  
8 204, and the shape/verb runtime system 220.

### 9 10 **System Component Functionality**

11 Typically, a designer relies on external modeling and animation systems  
12 and/or geometry and motion capture technology (external tools 208) to create  
13 initial forms and motions. In the context of this document, these initial forms and  
14 motions are referred to as example forms and example motions, or simply  
15 “examples”. As an illustration, consider an object that has a shape. One way of  
16 capturing the object’s shape might be to project structured light onto the object so  
17 that a camera and computer can record the object’s shape. This shape is digitized  
18 and can then be used in graphics applications. For a situation that involves  
19 motion, a person might be outfitted with a number of reflectors and then filmed so  
20 that a computer can capture their motion. In each of the these systems, once the  
21 images are captured they are very difficult to modify. In a traditional system,  
22 these captured images from externals tools 208 would flow directly to an  
23 application that would simply pull out the captured images and run them as  
24 appropriate in known ways.  
25

1 In the illustrated and described embodiment, shape and verb offline tools  
2 214 enable the designer to organize these example forms or motions that serve as  
3 input into shape and verb generator 204. To do so, the designer can choose a set  
4 of adjectives that characterize the forms or a set of adverbs that characterize the  
5 motions. The adjectives or adverbs define an abstract space and each adjective or  
6 adverb represents a separate axis in this abstract space. For instance, adjectives  
7 that describe the form of a human arm may include gender, age, and elbow bend.  
8 These adjectives are used to define the axes in an abstract space.

9 Consider Fig. 3, for example. There, each of the adjectives is seen to define  
10 an axis in the abstract space. The adjective "gender" can range in values from  
11 male to female; the adjective "age" can range in values from old to young; and the  
12 adjective "elbow bend" can range in value from no bend to full bend. These axes  
13 are of interest because the form of a human arm changes depending on whether it  
14 belongs to a male or female, or whether the person is old or young. The arm also  
15 deforms when the skeleton bends. In the latter case, deformation does not refer to  
16 the rigid body transformations induced by, for instance, bending the elbow, but  
17 rather the more subtle non-rigid changes in muscles and skin.

18 Adverbs for a walk may include the walker's mood and aspects of the  
19 environment in which the walk takes place. A happy walk is quite different from a  
20 sad or angry walk. Walks also differ with the slope or the surface walked on.

21 Once the designer has defined the abstract space using shape and verb  
22 offline tools 214 (Fig. 2), each example form or motion is annotated with its  
23 location in the abstract space. Annotation of an example's location in the abstract  
24 space is a mathematical characterization of the example's location. In addition,  
25 motion examples can be tagged with so-called keytimes, such as the moment when

1 each foot touches the ground. The keytimes provide a means to perform automatic  
2 time warping at runtime. Details of the time warping can be found in Rose et al,  
3 *Verbs and Adverbs: Multidimensional motion interpolation*, IEEE Computer  
4 Graphics and Applications 18, 5 (Sept. 1998), pps. 32-40.

5 Based on the annotated examples and the abstract space, shape and verb  
6 generator 204 solves for the coefficients of a smoothly-varying interpolation of the  
7 forms and motions across the abstract space. These coefficients provide the means  
8 to interpolate between the example forms and motions at runtime. In the context  
9 of this document, the output produced by the shape and verb generator 204 is  
10 referred to as a “shape” when interpolating forms, and as a “verb” when  
11 interpolating motions.

12 At runtime, application 206 chooses, at each moment in time, desired  
13 values for the adjectives or adverbs, thus defining a specific location in the  
14 abstract space. For instance, a character can be set to be happy or sad or anywhere  
15 in between; an arm can be specified to be more male or female, or to respond to  
16 the bending of the elbow. The shape and verb runtime system 220 then responds  
17 to the selected location in the abstract space by efficiently blending the annotated  
18 examples to produce an interpolated form or motion.

19 The number of adverbs or adjectives defines the dimension of the abstract  
20 space. In this document,  $D$  is used to denote the dimension. The number of  
21 examples included in a construction of a shape or verb are denoted by  $N$ . Each  
22 motion in a verb as well as each form in a shape is required to have the same  
23 structure. That is, all forms in a shape must have the same number of vertices with  
24 the same connectivity. Since example forms must all have the same topological  
25 structure, we do not need to address the correspondence problem here. Examples

must each define the same set of DOFs, but may do so in a variety of ways. Different curve representations, number of control points, etc., can all coexist within a single set of examples. All example motions in a particular verb must also represent the same action. A set of example walks, for instance, must all start out on the same foot, take the same number of steps, and have the same arm swing phase. Thus, all examples (forms or motions) must have the same number of degrees of freedom (DOFs). The number of DOFs, denoted  $M$ , equals three times the number of vertices in a form (for the  $x, y, z$  coordinates). In the case of a motion, the number of DOFs is the number of joint trajectories times the number of control points per curve.

In the discussion that follows, a number of different symbols are used to mathematically describe and illustrate the inventive approach. The table just below summarizes or explains these symbols.

Object	Variable	Subscript	Range
<i>Example</i>	$X$	$i$	$1 \dots N$
DOF	$x$	$i$	$1 \dots N$
		$j$	$1 \dots M$
Point in Adverb Space	$\mathbf{p}$	$i$	
Radial basis	$R$	$i$	
Radial Coefficient	$r$	$i, j$	
Linear Basis	$A$	$l$	$0 \dots D$
Linear Coefficient	$a$	$i, l$	
Distance	$d$	$i$	

Consider that an example is a collection of numbers that numerically describe the DOFs. For instance, a shape is composed of many vertices and triangles. Each vertex can be considered as three numbers and there may be a few hundred vertices that describe a shape. Typically, there may be a few thousand numbers that are associated with each example and that describes the example. With respect to motion, such is analogously described by a collection of numbers that might define, for example, the motion of an arm during a defined period of time. An example is denoted mathematically as  $X_i$ , where:

$$X_i = \{x_{ij}, \mathbf{p}_i, K_m : i = 1 \dots N, j = 1 \dots M, m = 0 \dots \text{NumKeyTimes}\}$$

Each  $x_{ij}$ , the  $j^{\text{th}}$  DOF for the  $i^{\text{th}}$  example represents a coordinate of a vertex or, in the case of a motion, a uniform cubic B-spline curve control point.  $\mathbf{p}_i$  is the location in the abstract space assigned to the example.  $K$  is the set of keytimes that describe the phrasing (relative timing of the structural elements) of the example in the case of a motion. Based on the keytime annotations for each motion example, the curves are all time-warped into a common generic time frame. Keytimes are, in essence, additional DOFs and are interpolated at runtime to undo the time-warp. See Rose et al., referenced above, for details on the use of keytimes.

### Shape and Verb Generation

Given a set of examples (i.e. shapes or motions), the challenge is to ascertain mathematically how a non-example shape or motion is going to look.

1 Recall that the non-example shape or motion constitutes a shape or motion that  
2 was not defined by a designer. In the illustrated and described embodiment, a set  
3 of weights is developed for a particular point in the abstract space. The particular  
4 point represents a point that is not an example and for which a shape or motion is  
5 to be derived. The set of weights for that point consists of a weight value for each  
6 example. The individual weight values are then blended together to define the  
7 shape or motion at the given point.

8 Fig. 4 is a flow diagram that describes steps in a method in accordance with  
9 the described embodiment. The method can be implemented in any suitable  
10 hardware, software, firmware, or combination thereof. In the illustrated example,  
11 the method is illustrated in software. Step 400 provides a point in an abstract  
12 space for which a form or motion is to be defined. The abstract space is populated  
13 with examples, as described above. Step 402 develops a set of weights or weight  
14 values for that particular point based upon the predefined examples that populate  
15 the abstract space. That is, for each example, a weight value is derived. Step 404  
16 then combines the individual weights and their corresponding examples to define a  
17 shape or motion for the point of interest. In the illustrated and described  
18 embodiment, each weight value is used to adjust all of the corresponding  
19 coefficients of an example that define that example. This constitutes an  
20 improvement over past methods, as will be described in more detail below.

21 That is, given a set of examples, a continuous interpolation over the abstract  
22 space is generated for the shape or verb. The goal is to produce, at any point  $\mathbf{p}$  in  
23 the abstract space, a new motion or form  $X(\mathbf{p})$  that is derived through interpolation  
24 of the examples. When  $\mathbf{p}$  is equal to the position  $\mathbf{p}_i$  for a particular example  $i$ , then  
25



1  $X(p)$  should equal  $X_i$ . In between the examples, smooth intuitive changes should  
2 take place.

3 In Rose et al., referenced above, each B-spline coefficient was treated as a  
4 separate interpolation problem (e.g. 1200 separate problems in their walk verb)  
5 which leads to inefficiencies. In the illustrated and described embodiment, a  
6 cardinal basis is developed where one basis function is associated with each  
7 example. As a cardinal basis, each basis function has a value of 1 at the example  
8 location in the abstract space and a value of 0 at all other example locations. This  
9 specification guarantees an exact interpolation. For each DOF, the bases are  
10 simply scaled by the DOF values and then summed.  
11 This approach is more efficient than the approach set forth in Rose et al. in the  
12 case of animations and, the approach can be applied to shape interpolation.

13 Fig. 5 is a flow diagram that describes steps in a method in accordance with  
14 the described embodiment. The steps can be implemented in any suitable  
15 hardware, software, firmware, or combination thereof. In the illustrated example,  
16 the steps are implemented in software. Step 500 provides a plurality of examples.  
17 These examples can be generated by designers as described above. Step 502  
18 develops a cardinal basis where one basis function is associated with each  
19 example. Step 504 then interpolates to find a point in abstract space that is a  
20 blended combination of each of the examples provided in step 500.

21 One specific example of how this can be done is described below. It will  
22 be appreciated that the specific example constitutes but one way of implementing  
23 the described embodiment. Other ways of implementing the described  
24 embodiment can be used without departing from the spirit and scope of the  
25 claimed subject matter.

## Selecting Cardinal Basis Function Shapes

In the illustrated example, a specific shape of the individual cardinal basis functions is selected. In the present situation, the problem is essentially one of scattered data interpolation, as we have few data points, the examples, in a relatively high dimensional space. Most published scattered data interpolation methods focus on one and two-dimensional problems. Linear interpolation using Delauney triangulation, for instance, does not scale well in high dimensions. Based on the present need to work in high dimensions with very sparse samples, an approach is adopted which is a combination of radial basis functions and low order (linear) polynomials.

For a given point in abstract space, the defined cardinal basis gives the weight values for each of the examples that are to be used to define the shape or motion for the point in abstract space. The specific cardinal basis that was used in experimentation has the form:

$$w_{i_1}(\mathbf{p}) = \sum_{i_2=1}^N r_{i_2, i_1} R_{i_2}(\mathbf{p}) + \sum_{l=0}^D a_{i_1, l} A_l(\mathbf{p})$$

Equation 1

where  $r_{i_2, i_1}$  and  $R_{i_2}$  are, respectively, the radial basis function weights and radial basis functions themselves, and the  $a_{i_1, l}$  and  $A_l$  are the linear coefficients and linear bases. The subscripts  $i_1$  and  $i_2$  both indicate example indices. Given these bases, the value of each DOF is computed at runtime based on the momentary location,  $\mathbf{p}$ , in the abstract space. The value of each DOF,  $x_j$ , at location  $(\mathbf{p})$  is, thus, given by:

$$x_j(\mathbf{p}) = \sum_{i=1}^N w_i(\mathbf{p}) x_{i,j}$$

Equation 2

The equation directly above constitutes a weighted sum of all of the examples for each DOF at a location  $\mathbf{p}$ . This weighted sum provides the shape or motion of the point in abstract space for which an example does not exist. In the past, methods have not calculated a single weight for each example. Rather, individual weights for each of the coefficients that describe a single example were calculated. Where there are many examples, this computational load is quite burdensome and results in a very slow evaluation that is not suitable for runtime animation. Here, however, a single weight is calculated for each example and then applied to the coefficients that define that example. The modified examples are then combined to provide the resultant shape or motion. This approach is advantageous in that it is much faster and is readily suited for runtime animation. Examples of experimental results are given below that demonstrate how the inventive systems and methods improve on past techniques.

The linear function in the above approach provides an overall approximation to the space defined by the examples, and permits extrapolation outside the convex hull of the locations of the examples, as will be understood by those of skill in the art. The radial bases locally adjust the solution to exactly interpolate the examples. The linear approximation and the radial bases are discussed below in more detail.

### Simple One Dimensional Example

Before exploring in detail the derivation of the cardinal basis, a simple one dimensional example is given to assist the reader in envisioning the derivation that is described below.

Consider Fig. 6 which is a graph of a function where the  $x$ -axis represents a continuum of all of the possible positions in abstract space and the  $y$  axis represents the value of a weight  $w_I$ . Consider also that there are five examples (Ex 1 – Ex 5) that have been defined such that, for any point in abstract space we wish to derive a set of weights ( $w_1$ –  $w_5$ ) such that a combination of the weights yields a blended new example that is a realistic representation of any point in the abstract space. The goal of the derivation of the cardinal basis is that the basis function or weight has a value of 1 at the example location and 0 at all other example locations.

Here, notice that the  $x$  axis is labeled with five points (Ex 1, Ex 2, Ex 3, Ex 4, and Ex 5). These points constitute the five defined example locations in the abstract space. Notice that the weight value  $w_I$  for each of these points is 1 for example 1 (Ex 1) and 0 for all of the other examples. A point of interest can be envisioned as sliding back and forth along the  $x$  axis and as it does so, its weight coefficient  $w_I$  will change in accordance with the function that is defined for  $w_I$ . With respect to this function, we know that it will have a value of 1 at the point that corresponds to Example 1 and will be 0 elsewhere. What we do not know, however, is the shape of the function at the places between the points that correspond to the examples. This is something that is derived using Equation 1 above.

Equation 1 is essentially an equation that represents a sum of two different parts—a radial basis part (the term on the left side of the equation) and a linear

1 part (the term on the right side of the equation). The weight function for an  
2 example at some point  $\mathbf{p}$  is equal to the summed weighted radial basis functions  
3 and the summed linear basis. The linear basis is essentially a plane equation for a  
4 plane in the abstract space. For a 4-dimensional space, the equation would be a  
5 hyperplane, 3-dimensional space, the equation would be a plane, for a 2-  
6 dimensional space, the equation would be a line.

7 In the discussion just below, the section entitled "Linear Approximation"  
8 describes characteristics of the linear portion of equation 1 and the section entitled  
9 "Radial Basis" describes characteristics of the radial basis portion of equation 1.

### 10 11 **Linear Approximation**

12 In the illustrated example, we first would like to form the best (in the least  
13 squares sense) linear approximation for each DOF based on the examples given  
14 for that DOF. In other words, we would like to find the hyperplane through the  
15 abstract space that comes closest to approximating the example values of that  
16 DOF. This defines  $M$  separate least squares problems, one for each DOF variable.  
17 However, since each example places all variables at a single location,  $\mathbf{p}$ , we can  
18 instead determine a set of  $N$  hyperplanes, one for each example, that form a basis  
19 for the  $M$  hyperplanes. The basis hyperplanes are derived by fitting a least squares  
20 hyperplane to the case where one example has a value of 1 and the rest have a  
21 value of 0.

22 Consider the following:

$$23 \quad \mathbf{p}_h \mathbf{a} = \mathbf{F},$$

24 where  $\mathbf{p}_h$  is a matrix of homogeneous points in the abstract space (i.e., each  
25 row is a point location followed by a 1),  $\mathbf{a}$  are the unknown linear coefficients, and

1 **F** is a *Fit* matrix expressing the values we would like the linear approximation to  
2 fit. In this case, **F** is simply the identity matrix since we are constructing a  
3 cardinal basis. Later, **F** will take on a slightly different structure as we discuss  
4 reparameterization of the abstract space (see “Reparameterization” section below).

5 Fig. 7 shows the linear approximation  $A_1$ , and three exemplary radial basis  
6 functions associated with three different examples for a simple one-dimensional  
7 abstract space. The three examples are located at  $p = 0.15, 0.30, 0.75$ . The straight  
8 line labeled  $A_1$  is the line that fits best through  $(0.15, 1), (0.30, 0), (0.75, 0)$ . A  
9 “best fit” hyperplane for any particular DOF could then be evaluated by simply  
10 scaling the bases by the DOF values and summing.

11 In the Fig. 7 example, the  $R_x$  terms represent the basic shape of the radial  
12 basis function. The lowercase “r” associated with each linear basis function  
13 represents an amount by which the basis function is scaled in the vertical direction.  
14 The width of the basis functions, in this example, is implied by the distance to the  
15 nearest other example. For instance,  $p_1$  and  $p_2$  are close to one another and  
16 constitute the closest examples relative to the other. In this particular example, the  
17 width of the radial parts of the function is four times the distance to the nearest  
18 example. Hence, the width of  $p_1$  and  $p_2$  are the same, only shifted relative to one  
19 another. The width for example  $p_3$  is much larger because it is further away from  
20 its nearest example ( $p_2$ ). Fig. 7 simply illustrates an example of equation 1 above.

## 21 22 **Radial Basis**

23 Radial basis interpolation is known and is discussed in Micchelli,  
24 *Interpolation of scattered data: Distance matrices and conditionally positive*  
25 *definite functions*, Constructive Approximation 2 (1986), and in a survey article

Powell, *Radial basis functions for multivariable interpolation: A review*, Algorithms for Approximation, J.C. Mason and M.G. Cox, Eds. Oxford University Press, Oxford UK, 1987, pps. 143-167. Radial basis functions have been used in computer graphics for image warping as described in Ruprecht et al., *Image warping with scattered data interpolation*, IEEE Computer Graphics And Applications 15, 2 (Mar. 1995), 37-43; Arad et al., *Image warping by radial basis functions: Applications to facial expressions*, Computer Vision, Graphics, and Image Processing 56, 2 (Mar. 1994), 161-172; and for 3D interpolation as described in Turk et al., *Shape transformation using variational implicit functions*, Computer Graphics (Aug. 1999), pps. 335-342 (Proceedings of SIGGRAPH 1999).

Given the linear approximation described above, there still remain residuals between the example values  $x_{ij}$  and the scaled and summed hyperplanes. To account for the residuals, one solution might be to associate a radial basis with each DOF, as described in Rose et al mentioned above. Instead, we proceed as before and use the radial bases to account for the residuals in the cardinal bases. The residuals in the cardinal bases are given by:

$$q_{i1i2} = \delta_{i1i2} - \sum_{l=1}^D a_{i2l} A_l(\mathbf{p}_{i1}), \text{ where } \delta_{ij} = \begin{cases} 1 & \text{when } i=j \\ 0 & \text{otherwise} \end{cases}$$

To account for these residuals, we associate  $N$  radial basis functions with each example. Since there are  $N$  examples, we, thus, need  $N^2$  radial basis functions. We solve for the weights of each radial bases,  $r_{i2i1}$ , to account for the residuals such that, when the weighted radial bases are summed with the hyperplanes, they complete the cardinal bases.

1 This leaves the issue of choosing the specific shape of the radial bases and  
2 determining the radial coefficients. Radial basis functions have the form:

$$R_i(d_i(\mathbf{p}))$$

3  
4  
5 where  $R_i$  is the radial basis associated with  $X_i$  and  $d_i(\mathbf{p})$  is a measure of the  
6 distance between  $\mathbf{p}$  and  $\mathbf{p}_i$ , most often the Euclidean norm  $\|\mathbf{p} - \mathbf{p}_i\|$ . There are a  
7 number of choices for this basis. Rose et al., referenced above, chose a basis with  
8 a cross section of a cubic B-spline centered on the example, and with radius twice  
9 the Euclidean distance to the nearest other example. Turk et al., referenced above,  
10 selected the basis:

$$R(d) = d^2 \log(|d|)$$

11  
12  
13  
14 as this generates the smoothest (thin plate) interpolations. In the described  
15 embodiment, both bases were tried. Bases with the B-spline cross-section were  
16 found to have compact support, thus, outside their support, the cardinal bases are  
17 simply equal to the linear approximation. They therefore extrapolate better than  
18 the smoother, but not compact, radial basis functions used by Turk et al.  
19 Accordingly, we chose to use the B-spline bases for this extrapolation property.

20 The radial bases weights,  $r_{i_2 i_1}$ , can now be found by solving the matrix  
21 system:

$$\mathbf{Q}\mathbf{r} = \mathbf{q}$$

22  
23  
24 where  $\mathbf{r}$  is an  $N \times N$  matrix of the unknown radial bases weights and is  
25 defined by the radial bases such that,  $Q_{i_2 i_1} = R_{i_2}(p_{i_1})$  the value of the unscaled radial



1 basis function centered on example  $i_2$  at the location of example  $i_1$ . The diagonal  
2 terms are all  $2/3$  since this is the value of the generic cubic B-spline at its center.  
3 Many of the off diagonal terms are zero since the B-spline cross-sections drop to  
4 zero at twice the distance to the nearest example.

5 Referring back to Fig. 7, three radial basis functions are seen to be  
6 associated with the first of the three examples. If these three radial basis functions  
7 are summed with the linear approximation, we get the first cardinal basis, the line  
8 labeled  $w_1$  in Fig. 8. Note that the first cardinal basis passes through 1 at the  
9 location of the first example, and is 0 at the other example locations. The same is  
10 true for the other two cardinal bases  $w_2$  and  $w_3$ .

11 Given the solutions for the radial basis weights, we now have all the values  
12 needed to evaluate equations 1 and 2 at runtime. The line designated  $\sum wx$  in Fig.  
13 8 is simply the three cardinal bases scaled by the example values and summed as  
14 in equation 2.

## 15 16 **Shape and Verb Runtime System**

17 At runtime, an application selects a point of interest in the abstract space.  
18 This point may move continuously from moment to moment, if for instance, a  
19 character's mood changes or the character begins to walk up or downhill. The  
20 shape and verb runtime system 220 (Fig. 2) takes this point and generates an  
21 interpolated form or motion and delivers it to the application for display.

## 22 23 **Complexity Issues**

24 The runtime complexity of the formulation in equations 1 and 2 is  
25

$$MN + N(N + S) = MN + N^2 + NS$$

where  $M$  is the number of DOF variables,  $N$  is the number of examples, and  $S$  is the dimension plus one of the abstract space. In Rose et al., referenced above, there was a separate linear hyperplane and radial basis per DOF (approximately 1200 B-spline coefficients in their walking verb). The complexity in their formulation is given by:

$$M(N + S) = MN + MS$$

In our formulation above there is only one set of linear bases and one set of radial bases per example. The efficiency is gained due to the fact that there are many fewer examples than DOF variables. For instance, if  $M = 1200$ ,  $N = 10$ , and  $S = 6$ , then the comparison is  $MN + N^2 + NS = 12,160$  vs.  $MN + MS + N^2 = 19,200$ . Due to some additional constants, the efficiencies are actually a bit better. We have implemented both Rose et al.'s formulation and the one presented here and found an approximate increase in speed at a factor of about 2 in most cases.

### **Reparameterization and Modification of Shapes and Verbs**

Once a shape or verb has been generated, the abstract space can be quickly explored to see the interpolated forms and motions. It is possible that some of the regions in this space are not entirely to the designer's liking. For instance, Fig. 9 shows a 2D space of interpolations of a simple three-dimensional shape. Blue forms indicate examples and grey forms are the result of the shape and verb runtime system 220 (Fig. 2).

Problem regions can be seen in each of the two upper corners. Both corners exhibit cracks due to surface interpenetration. Another type of problem

1 can occur when an interpolation changes more quickly than desired in one region  
2 and more slowly than desired in a neighboring region. For instance, at the point in  
3 the abstract space that is half way between a happy walk and a sad walk, the  
4 character may appear to be more sad than happy, rather than neutral.

5 Two methods are typically available to modify the results produced by the  
6 shape and verb generator 204 (Fig. 2). In one method, a new example can be  
7 inserted at the problem location and a new verb or shape generated. A quick way  
8 to bootstrap this process is to use the undesirable results at the problem location as  
9 a "starting point", fix up the offending vertices or motion details, and reinsert this  
10 into the abstract space as a new example. In the second method, the designer first  
11 selects an acceptable interpolated form or motion from a location in the abstract  
12 space near the problem region. Next, the designer moves this form to a location in  
13 the problem region. This relocated form will be treated exactly the same as an  
14 example. This relocated form is referred to as a "pseudo-example" in this  
15 document.

16 Going back to the problem with the walk verb described above, a neutral  
17 walk cycle found elsewhere in the abstract space may be moved to the half way  
18 point between happy and sad walks.

19 The place from which the pseudo-example is taken will be point  $p$ . We  
20 denote the new position to which it is moved to as  $\tilde{p}$ . The pseudo-example is not  
21 a "true" example as it is a linear sum of other examples. However, it acts like an  
22 example in terms of changing the shape of the radial basis functions. In essence,  
23 this method reparameterizes the abstract space. Reparameterization is very  
24 effective in producing shapes of linked figures, as we demonstrate later with a  
25 human arm shape.

Fig. 10 shows this technique applied to the 2D space of 3D forms. The space has been reparameterized with two pseudo-examples shown in red. Arrows indicate the locations from which the pseudo-examples were drawn.

Creating the reparameterized verb or shape proceeds much as before. We introduce a few new notations, however, before we formalize the reparameterization approach. We have already mentioned  $\tilde{\mathbf{p}}$ , the new location of the pseudo-examples.  $\tilde{\mathbf{p}}$  also includes the real example locations during reparameterization. We, furthermore, use a tilde to denote the new radial basis weights,  $\tilde{r}$ , the new linear coefficients  $\tilde{a}$ , and the new cardinal bases  $\tilde{w}$ . We also denote the total number of real and pseudo-examples as  $\tilde{N}$ .

As before, the form or motion at any point in the abstract space is:

$$x_j(\mathbf{p}) = \sum_{i_1=1}^{\tilde{N}} \tilde{w}_{i_1}(\mathbf{p}) x_{i_1 j}$$

Equation 3

Note that the interpolated form or motion still only includes a sum over the real examples. In other words, there are still only  $N$  cardinal bases after the reparameterization. The pseudo-examples reshape these  $N$  bases from  $w$  to  $\tilde{w}$ .

$$\tilde{w}_{i_1}(\mathbf{p}) = \sum_{i_2=1}^{\tilde{N}} \tilde{r}_{i_2 i_1} R_{i_2}(\mathbf{p}) + \sum_{l=0}^D \tilde{a}_{i_1 l} A_l(\mathbf{p})$$

Equation 4

Also as before  $\tilde{\mathbf{p}}_h \tilde{\mathbf{a}} = \tilde{\mathbf{F}}$ . However,  $\tilde{\mathbf{F}}$  is no longer an identity matrix. It now has  $\tilde{N}$  rows and  $N$  columns. Assuming the new pseudo-examples are all located at the bottom of the matrix, the top  $N \times N$  square is still an identity. The

lower  $\tilde{N}-N$  rows are the values of the original cardinal bases  $w$  at the location where the pseudo-examples were taken from (see equation 4). That is, in each row, the *Fit* matrix contains the desired values of the  $N$  cardinal bases, now at  $\tilde{N}$  locations. These are 1 or 0 for the real example locations and the original cardinal weights for the pseudo-examples.

The radial portion of the cardinal basis construction proceeds similarly. The residuals are now:

$$\tilde{q}_{i_1 i_2} = \tilde{F}_{i_1 i_2} - \sum_{l=0}^D \tilde{a}_{i_2 l} A_l(\mathbf{p}_{i_1})$$

Note that instead of the Kronecker delta, we now have the values from  $\tilde{F}$ .

The coefficients,  $\tilde{r}_{i_2 i_1}$ , are now found by solving the matrix system,

$$Q\tilde{r} = \tilde{q}$$

As before,  $Q_{i_1 i_2}$  has terms equal to  $R_{i_2}(\tilde{p}_{i_1})$ , however, these terms now include the new pseudo-example locations. As a consequence, the radii of the radial basis functions may change.

Fig. 11 is a flow diagram that describes steps in a reparameterization method in accordance with the described embodiment. Step 1100 defines a set of examples that pertain to a form or motion that is to be animated. The examples are provided relative to a multi-dimensional abstract space. In the illustrated example, this step can be implemented by a animation designer using a set of software tools to design, for instance, example objects or character motions. Step 1102 examines a plurality of forms or motions that are animated with the abstract space from the defined set of examples. In the illustrated and described

embodiment, the forms or motions can be animated from the examples using the techniques described above. Step 1104 then identifies at least one form or motion that is undesirable. This step can be manually performed, e.g. by the designer physically examining the animated forms or motions, or it can be performed by software that looks for undesirable forms or motions. An example of an undesirable form or motion is one that, for example, exhibits undesirable characteristics such as surface interpenetration due to interpolation issues. Once the undesirable forms or motions are identified, step 1106 selects a form or motion from a location within the abstract space that is proximate a location that corresponds to the undesirable form or motion. This step can be implemented by a designer actually physically selecting the form or motion. Alternately, it can be performed by software. Once the form or motion is selected, step 1108 replaces the undesirable form or motion with the selected form or motion to provide a pseudo-example that constitutes a linear sum of the original examples. Processing as far as subsequent rendering operations can take place as described above.

### **Shapes and Skeletons**

Animated characters are often represented as an articulated skeleton of links connected by joints. Geometry is associated with each link and moves with the link as the joints are rotated.

If the geometry associated with the links is represented as rigid bodies, these parts of the character will separate and interpenetrate when a joint is rotated. A standard way to create a continuous skinned model is to smoothly blend the joint transforms associated with each link of the character. This is supported in

1 current graphics hardware, making it an attractive technique. This simple method  
2 exhibits some problems, however.

3 We use an arm bending at the elbow to discuss the problems associated  
4 with transform blending and to demonstrate how we can overcome these  
5 difficulties in the context of shape blending.

6 To perform the simple transform blending, a blending weight  $\alpha$  is assigned  
7 to each vertex. For instance, for the elbow bending, vertices sufficiently below the  
8 elbow would have weight  $\alpha = 1$ , and those sufficiently above,  $\alpha = 0$ . Those  
9 vertices near the elbow would have  $\alpha$  values between 0 and 1, as will be  
10 understood by those of skill in the art.

11 We denote the transformation matrix for the upper arm as  $T_0$ , and the  
12 transformation for the lower arm as  $T_1$ . We use superscripts to denote the amount  
13 of rotation of the joint. Thus, as the elbow rotates, we say that  $T_1$  changes from  
14  $T_1^0$  when the elbow is straight to  $T_1^1$  at a bent position. In between, the transform  
15 is  $T_1^\beta$  for a bending amount  $\beta$ . We refer to the position of the arm when  $T_1 = T_1^0$   
16 as the rest position.

17 We denote the position of a vertex as  $x_0$  when the transformation of the  
18 joint is  $T_1^0$ . The simple transform blending is defined as:

$$x = \alpha T_1^\beta x_0 + (1-\alpha) T_0 x_0$$

19  
20  
21 where  $\alpha$  is the blending weight assigned to that vertex. Unfortunately,  
22 linearly summed transformation matrices do not behave as one would like. The  
23 result is that the skin appears to shrink near the rotating joint (see Fig. 12). In  
24 addition, simple transform blending does not provide a rich set of tools for  
25 creating effects such as a muscle bulging as the elbow bends.

In the illustrated and described embodiment, we provide solutions to both problems through the use of our inventive shape blending techniques. To do this, the designer first creates examples of a straight arm,  $X_{\beta=0}$ , and a bent arm,  $X_{\beta=1}$ , for, say, a 90-degree bend (Fig. 11). The bent arm includes any muscle bulging or other deformations the designer wants to include. We obviously cannot simply perform a direct geometric blend between these forms because at 45 degrees, the lower arm would have shrunk considerably, since the chord from the fingertip of the bent arm to the fingertip of the straight arm passes nearer to the elbow than an arc would, as will be appreciated by those of skill in the art.

Instead, what we would like is a combination of transform blending and shape blending. Accordingly, we call the vertex on the bent arm, corresponding to  $x_0$  on the arm in the rest position,  $x^1$ . As in the simple blended transforms, we also have a blending weight,  $\alpha$ , associated with each vertex. We now seek a second arm in the rest position with a corresponding vertex,  $x_0^1$ , such that when it is subjected to the simple transform blending at an angle of  $\beta=1$  it will exactly match the bent arm specified by the designer. Thus,

$$x^1 = \alpha T_1^{\beta=1} x_0^1 + (1 - \alpha) T_0 x_0^1$$

Now we can solve for the vertices of this new arm in the rest position (see Fig. 13):

$$x_0^1 = (\alpha T_1^{\beta=1} + (1 - \alpha) T_0)^{-1} x^1$$

Equation 5



1 We can now first perform a geometric blend of the new arm in the rest  
2 position with the original arm in the rest position and then transform it. The result  
3 is:

$$x_{\beta} = (\alpha T_1^{\beta} + (1 - \alpha)T_0)(\beta x_0^1 + (1 - \beta)x_0)$$

4  
5  
6 Equation 6  
7

8 This geometric blend followed by the blended transform will match the  
9 original arm geometry when  $\beta=0$  and will match the bent arm created by the  
10 designer when  $\beta=1$ .

11 The arm model discussed above contains only two example forms, the arm  
12 in the rest position and the bent arm. It will be appreciated, however, that all of  
13 the techniques discussed above can be used to perform a multi-way blend of the  
14 articulated geometry. First, all example forms are untransformed to the rest  
15 position via equation 5. Then the multi-way shape blending is applied to the  
16 untransformed forms. Finally, the blending result is transformed in the normal  
17 way. In other words, the interpolation of the forms in the rest position simply  
18 replaces the latter half of equation 6.

19 Fig. 14 shows the straight arm example one, of 6 examples in this shape.  
20 The results above the straight arm are the result of naively blending transforms on  
21 this one example. On the bottom right is a blend of the 6 examples untransformed  
22 to the rest position. Finally, this strange blended form is pushed through the  
23 blended transformation to give the result shown in the upper right.

24 Fig. 15 is a flow diagram that describes step in a animation method in  
25 accordance with the described embodiment. Step 1500 defines examples forms in

1 different positions. In the illustrated, a first example form is defined in a first  
2 position and a second example form is defined in a second position. Step 1502  
3 then computes a form in the first position such that when the computed form is  
4 subjected to a transform blending operation that places the computed form in the  
5 second position, it will match the second example form. Step 1504 geometrically  
6 blends the computed form in the first position with the first example form in the  
7 first position to provide a geometrically blended form in the first position.  
8 Specific examples of this are given above. Step 1506 then transform blends the  
9 geometrically blended form to provide a form that matches the second example  
10 form in the second position.

## 11 **Results**

12 We have applied the paradigm described above to a number of shapes and  
13 linked-figure animations. We will discuss two of each. All performance statistics  
14 measure raw interpolation speeds and do not include time to render. All timings  
15 were gathered on a 450 Mhz Pentium II, Dell Precision 610- machine.  
16

## 17 **Simple Shapes**

18 The shape demonstrated in Fig. 10 includes 5 example forms, each  
19 consisting of 642 vertices and 1280 faces. Blending in this space can be done at  
20 5500 frames per second (fps). One step of Loop subdivision is applied to this  
21 mesh and the limit positions and normals are computed. This results in a mesh  
22 with 2562 vertices and 5120 faces that is rendered at runtime. A two dimensional  
23 abstract space was created with the horizontal axis being "bend" and the vertical  
24 being "thickness." The initial space had regions in the upper right and left corners  
25

1 where the blended surface was interpenetrating. This space was reparameterized  
2 by adding pseudo-examples near these locations that were not interpenetrating.  
3 When real examples were added instead, the blending slowed down to around  
4 4000 fps.

## 5 6 **Arm**

7 The arm was created by modifying Viewpoint models inside 3D-  
8 Studio/Max. The arm example has an underlying skeleton that has 3 rotational  
9 degrees of freedom - one for the shoulder, elbow and wrist. We have a fourth  
10 variable for the parameterized arm, gender. The abstract space has 8 real  
11 examples (straight arm, shoulder up, shoulder down and elbow bent for male and  
12 female models), and 6 pseudo-examples which are mostly placed to smooth out  
13 bulges where the shrinking from blended transforms was being overcompensated.  
14 These examples have 1335 vertices and 2608 faces each. This dataset can be  
15 interpolated at 2074 fps.

16 Fig. 16 is a visualization of a 2D space of arms with 4 real examples and 1  
17 pseudo-example. Elbow bend is parameterized along the horizontal axis and  
18 gender along the vertical. Extrapolation can produce some interesting results, as  
19 are indicated.

## 20 21 **Animation**

22 We constructed parameterized motions, verbs, for two different characters:  
23 a human figure named Stan and a robot warrior. While both bipedal, each moves  
24 in very different ways.  
25

1 Stan moves in a humanlike, though often exaggerated, manner, while the  
2 robot warrior moves very mechanically. Our technique had no trouble  
3 encompassing both forms. Our motion examples were chosen to encompass  
4 emotional states such as mopiness, cockiness, fear, anger, and for physical  
5 parameters such as damage and surface slope.

6 The verbs shown in Figs. 17-19 were constructed from a mix of motion  
7 capture and hand animated source. The motion capture data was gathered  
8 optically. The hand-animated source was generated on two popular software  
9 animation packages: 3D-Studio/Max for the robot, and Maya for some of Stan's  
10 data. This data is applied to a skeleton positioned with a translation at the root and  
11 a hierarchy of Euler angle joints internally. A technique similar to one discussed  
12 in Bodenheimer, et. al., *The process of motion capture: Dealing with the data*,  
13 Proceedings of the Eurographics Workshop on Computer Animation and  
14 Simulation (Sept. 1997), pps. 3-18, ensured that the angles remained well-  
15 behaved.

16 Our robot warrior has 60 degrees of freedom. Our robot data exhibits two  
17 primary kinds of variation: slope of walking surface and level of damage to the  
18 robot. Damage results in limping and slope causes the character to rotate the  
19 joints in the leg, foot, and waist, and shift weight to accommodate.

20 Fig. 17 shows a sampling of the robot walking verb. Green boxed figures  
21 are examples. The others are interpolations and extrapolations. The vertical axis  
22 shows robot health and the horizontal axis surface slope. Fig. 19 shows the robots  
23 in the center column of Fig. 17 rotated to the side to better show the character  
24 limping. We constructed a number of robot verbs: powered-down idle, power-up,  
25 powered-up idle, power-down, idle-to-walk, walk, and walk-to-idle. These verbs,

1 together with a transitioning mechanism can be combined to form a verb graph  
2 such as discussed in Rose, et. al., referenced above. The robot verbs evaluate at  
3 9700 fps in the absence of rendering. Thus, a robot evaluated using verbs can be  
4 run at 30 fps with only a 0.3% percent processor budget.

5 We also constructed a walking verb from a human figure, Stan, with 138  
6 degrees of freedom. We chose 3 emotional axes for this verb which seemed to  
7 encompass the gamut of variation present in the data. These were happy-to-sad,  
8 knowledgeable-to-clueless, and low-to-high-energy. Our example motions exhibit  
9 these subjective characteristics: easy-going, mopey, cocky, goofy, angry, afraid,  
10 and neutral. Fig. 18 shows a sampling of our walk along the knowledge and  
11 happiness axes. Unlike the robot data, which was designed to fall at regular places  
12 in the adverb plane, our Stan motions are less regular. As there are no restrictions  
13 upon example placement, we are able to construct verbs with irregular example  
14 placement, as shown by the green-boxed figures. Evaluating a pose for Stan is  
15 extremely efficient: 3900 fps or a 0.7% processor budget at 30 fps.

## 16 Conclusions

17 Shape and motion interpolation has been shown to be an effective and  
18 highly efficient way of altering shape and motion for real-time applications such  
19 as computer games. The front-end solution process is also efficient, (a fraction of  
20 a second), so a tight coupling between artist, solver, and interpolator provides a  
21 new way for an artist to work without fundamentally altering their workflow.  
22 With our system, artists can more easily extend their work into the interactive  
23 domain.  
24  
25

1 In the context of skeleton based figures, we are able to combine both shape  
2 blending with blended transforms to create a smoothly skinned character. We are  
3 able to overcome the limitations of blended transforms, while including the artist's  
4 input for how muscles deform as the skeleton moves, and still maintain interactive  
5 performance.

6 Although the invention has been described in language specific to structural  
7 features and/or methodological steps, it is to be understood that the invention  
8 defined in the appended claims is not necessarily limited to the specific features or  
9 steps described. Rather, the specific features and steps are disclosed as preferred  
10 forms of implementing the claimed invention.